

NOAA Cloud Evaluation

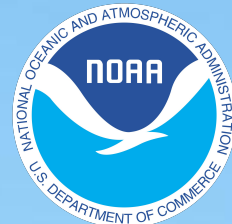
November 21, 2019

**Carolyn Pasti, President
RedLine Performance Solutions**

cpasti@redlineperf.com



GENERAL DYNAMICS
Information Technology



Acknowledgements

NOAA RDHPCS Sponsors

- Frank Indiviglio and Bernie Siebers

GDIT – NOAA RDHPCS Prime Contractor and Project Lead

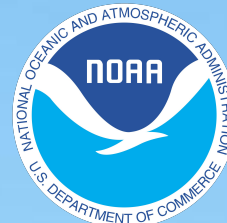
- Bill Asbury, Alan Powers, Rick Valencia, Raghu Reddy, Nathan Dauchy

RedLine Performance Solutions – GDIT Subcontractor

- Don Avart, Keith Ball, Andrew Qualkenbush, James Abeles

Cloud Service Providers (CSPs)

- Amazon, Microsoft, Penguin Computing, Google, Rescale
- All resources provided at no cost to the project through GDIT Alliance relationship or in specific support to the evaluation project



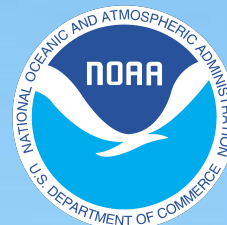
Overview

Goal

- Evaluate commercially available cloud platforms to support future mission functions of numerical weather prediction and support NOAA in defining the strategic goals of the HPCC program

Approach

- Establish baselines for FV3GFS and UPP benchmarks on NOAA systems
- Estimate deployment size to achieve 7 minutes per forecast day for FV3GFS
- Deploy virtual resources, install and run benchmarks
- Develop Singularity containers
- Develop cost model
- Evaluate against study objectives



Cloud Evaluation Findings

Objective	Findings
Performance	<ul style="list-style-type: none">• HPC cloud capabilities include exclusive node use, low-latency interconnects, node/traffic affinity and colocation features• High-performance network and HPC parallel file system are critical features• Run-time performance competitive with NOAA on-premise systems• Varying availability of low-latency interconnect between compute and storage
Cloud Management	<ul style="list-style-type: none">• Key ability is an HPC parallel file system “as a service”, which can support compute needs• Data storage and data transfer (ingress/egress) costs must be considered• Tiered storage: sync HPC parallel file system to lower-cost storage• Some CSPs have mature single-command deployment of compute, networking, and storage (after initial definition) – processes and ease of use varied• Cloud “on-demand” services do not include built-in method to limit usage by cost or hours used

Cloud Evaluation Findings (cont'd)

Objective	Findings
User Experience	<ul style="list-style-type: none">• Scheduler interfaces generally consistent/similar to on-premise HPC clusters• Utilized SSH to a login node (as with traditional clusters); APIs or Web access available in most cases• Data staging needs more efficient methods than just SCP; ease of use varied
Application Containers	<ul style="list-style-type: none">• Singularity framework is well suited to HPC (permissions model, no daemons, works with workload managers)• Need to make decisions regarding what goes into the container<ul style="list-style-type: none">• Application binaries and supporting libraries, environment variables• Some fixed input files may make sense• Containerized MPI applications introduce dependencies on MPI stack outside the container• Small to negligible overhead, straightforward to build

Further Considerations for Refining Cloud Strategy

- Determine best “cloud consumption model”
- Handling data: what should/can live in the cloud?
- Investigate lower cost nodes/processors (for storage and compute)
- Identify workflows capable of utilizing heavily discounted node instances
- Data staging and moving data to/from less expensive storage between job runs to reduce storage costs
- Explore performance of the other processors besides Intel which could reduce runtimes
- Evaluate Cloud bursting from on-premise resources
- Use Spack in a container as a portable build environment
- Establish cloud pilot for selected workflow(s)

